

WHAT IS CLAIMED IS:

1. A method for enabling a client to access a resource on a wide area network environment, comprising:

- (a) receiving a request for the resource from a connection that is associated with the client;
- (b) selecting a method for load balancing each request from the connection;
- (c) employing the selected method to select one of a plurality of server array controllers to handle each request from the connection to the resource; and
- (d) persistently referring each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the connection, wherein the selected server array controller associates the client's connection with the requested resource.

2. The method of Claim 1, wherein the client is a local domain name system (LDNS) server.

3. The method of Claim 2, wherein the selected load balancing method employs modulus arithmetic to select a virtual server managed by the selected server array controller, including:

- (a) converting the LDNS server's IP address into a first value;
- (b) dividing the first value by the total amount (N) of virtual servers that are available, each available virtual server corresponding to a particular value of N;
- (c) adding one to a remainder of the dividing action to create a second value; and
- (d) selecting the virtual server that corresponds to the particular value of N that is one greater than the second value.

4. The method of Claim 3, further comprising when the selected virtual server is determined to be down, the selected load balancing method performs further actions, including:

- (a) marking the down selected virtual server as unavailable;
- (b) decrementing the total amount of available virtual servers by 1 to create a new value for N; and
- (c) employing the new value for N to select another virtual server array.

5. The method of Claim 1, wherein the client employs a proxy server to send the request for the resource.

6. The method of Claim 1, wherein referring each request associated with the connection to the selected server array controller further comprises sending to the client an IP address associated with the selected server array controller.

7. The method of Claim 1, wherein the selected server array controller associates the client with a virtual server that includes the requested resource, wherein the virtual server is managed by the selected server array controller.

8. The method of Claim 1, further comprising storing information associated with the request.

9. The method of claim 9, wherein the information includes an IP address associated with the client, another IP address associated with the server array controller, a time stamp, and a time to live.

10. The method of Claim 1, wherein each request associated with the connection is made by any one of a plurality of LDNS servers.

11. The method of claim 1, wherein the selected server array controller manages at least one virtual server that provides access to the requested resource.

12. The method of Claim 1, wherein the selected method for load balancing includes one of static stripe, dynamic, and topological methods.

13. A method for enabling a client to access a resource on a wide area network environment, comprising:

(a) employing an extended domain name system (EDNS) server to load balance a request from a connection associated with the client for handling by a selected server array controller, each subsequent and timely request from the connection is persistently referred by the EDNS server for handling by the selected server array controller; and

(b) when the subsequent request is untimely, employing the EDNS server to load balance the request from the connection for handling by another selected server array controller.

14. The method of Claim 13, wherein the EDNS server includes a data store for storing information associated with the request.

15. The method of Claim 13, further comprising providing metrics from each server array controller to the EDNS server, wherein the metrics are used in the selection of the server array controller to handle the request from the connection.

16. The method of Claim 15, wherein the metrics include availability and round trip time.

17. The method of claim 14, further comprising another EDNS server that includes another data store for storing information associated with the request, wherein the EDNS and the other EDNS share request information between the data store and the other data store.

18. A modulated data signal including computer executable instructions, comprising:

(a) a receiver that receives a request for resources from a connection that is associated with the client;

(b) a selector that selects a method for load balancing each request from the connection and employs the selected method to select a server array controller to handle each request from the connection;

(c) a sender that refers each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the load balanced connection, wherein persistent communication is established between the client and the server.

19. An apparatus for providing persistent communication between a client and a server in a wide area network environment, comprising:

(a) means for receiving a request for resources from a connection that is associated with the client;

(b) means for selecting a method for load balancing each request from the connection;

(c) means for employing the selected method to select a server array controller to handle each request from the connection; and

(d) means for persistently referring each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the connection.

20. A computer data signal embodied in a carrier wave and representing computer executable instructions for providing persistent communication between a client and a server in a wide area network environment, comprising:

- (a) receiving a request for resources from a connection that is associated with the client;
- (b) selecting a method for load balancing each request from the connection;
- (c) employing the selected method to select a server array controller to handle each request from the connection; and
- (d) persistently referring each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the connection.

21. A computer-readable medium storing instructions causing a computer program to execute a computer process for providing persistent communication between a client and a server in a wide area network environment, comprising:

- (a) receiving a request for resources from a connection that is associated with the client;
- (b) selecting a method for load balancing each request from the connection;
- (c) employing the selected method to select a server array controller to handle each request from the connection; and
- (d) persistently referring each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the connection.

22. A method for providing persistent communication between a client and a server in a wide area network environment, comprising:

(a) means for receiving a request for resources from a connection that is associated with the client;

(b) means for selecting a method for load balancing each request from the connection;

(c) means for employing the selected method to select a server array controller to handle each request from the connection; and

(d) means for persistently referring each request associated with the connection to the selected server array controller until another method is selected for load balancing each request from the connection.